

LSI 01-021

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR PATENT

ON

*CYCLIC REDUNDANCY CHECKING FOR MANAGING THE COHERENCY OF
MIRRORED STORAGE VOLUMES*

BY

WILLIAM DELANEY
14715 SIEFKES STREET
WICHITA, KS 67230
CITIZEN OF USA

CERTIFICATE OF MAILING BY "EXPRESS MAIL"

"Express Mail" Mailing Label Number: EV 013 245 364 US

Date of Deposit: October 22, 2001

I hereby certify that this correspondence is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 C.F.R. § 1.10 on the date indicated above and is addressed to Box Patent Application, Assistant Commissioner for Patents, Washington, D.C. 20231

BY: Penny L. Flint
Penny L. Flint

[0001] *CYCLIC REDUNDANCY CHECKING FOR MANAGING THE COHERENCY OF
MIRRORED STORAGE VOLUMES*

FIELD OF THE INVENTION

[0002] The present invention generally relates to the field of mirrored data storage and particularly to a protocol and controller suitable for controlling the coherency of mirrored data storage volumes.

BACKGROUND OF THE INVENTION

[0003] Currently, information handling systems are often integrated into groups with multiple users accessing and storing data. For various reasons, groups of information handling system users may be located in separate areas such as in different offices, or in different geographic locations. Even with multiple users located in separate locations the need may arise to generate a unified data storage system. For example a common database of information may be desirable, such as a customer database for a business with different geographic locations. As a result, different storage volumes may need to be maintained and integrated to store and manage data.

[0004] When common volumes of information are stored a back-up of data may be needed such as by using other data storage volumes mirroring the primary storage volume to protect the data in the case of natural disasters, data corruption or the like. Additionally, the separate user locations may need to continue to function should a disruption in communication occur between the various locations.

[0005] Should a disruption in communication occur, the storage volumes may become incoherent, and as a result a secondary, mirroring storage volume may no longer represent the data on the primary storage volume. If the primary and the secondary storage volumes no longer mirror each other the need to reestablish coherency would require the communication of the data stored on the primary volume to the secondary

volume. A transmission of the primary volume to the secondary volume is undesirous due to the typically large volume of information and the time required to transmit.

[0006] Presently, to avoid the need to communicate the contents of the primary storage volume in the event of a break in communication, a bitmap is typically maintained to track the changes in the primary volume while communication between the primary and the secondary volumes is disrupted. A bitmap is utilized to provide a list of changes occurring to the primary storage volume so that the data written to the primary volume may be updated to the secondary volume when communication is reestablished. Thus, the need to transfer the entire contents of the primary storage volume to the secondary storage volume upon reestablishment of communication is eliminated. Generally, bitmaps of mirrored storage volumes are maintained in non-volatile memory to ensure their protection in the event of a power failure. Additionally, bitmaps are maintained in non-volatile memory because the mirrored storage volume system had no other method for establishing coherency besides transmitting the contents of the primary storage volume to the mirroring storage volume. Thus resulting in an input/output intensive operation.

[0007] Therefore, it would be desirable to provide a mirrored storage volume system capable of eliminating the need for a costly, time consuming non-volatile stored bitmap system and method for ensuring the coherency of mirrored volumes or in the alternative a mirrored volume system in which a bitmap may be maintained in volatile memory.

SUMMARY OF THE INVENTION

[0008] Accordingly, the present invention is directed to a mirrored storage volume system and methods for managing a mirrored storage volumes and restoring coherency in mirrored storage volumes. In a first aspect of the present invention a mirrored storage volume system includes a primary storage controller, a secondary storage controller, a primary storage volume, a secondary storage volume and a communication channel

linking the primary controller to the secondary controller. The primary and secondary controllers may be capable of conducting cyclic redundancy checking scans on their respective storage volumes. Furthermore, the CRC scans of the primary and secondary storage volumes may be compared to determine differences between the primary and the secondary storage volumes so that the secondary storage volume may be updated with data from the primary storage volume to ensure the secondary storage volume mirrors the primary storage volume.

[0009] In a further example a volatile memory may be linked to the primary controller so that if communication between the primary storage controller and the secondary storage controller is interrupted, a coarse grain bitmap may be maintained to update changes occurring on the primary storage volume to the secondary storage volume once communication is reestablished.

[0010] In a second aspect of the present invention, a method for restoring coherency in mirrored storage volumes is disclosed. The method includes conducting a CRC scan of a primary and secondary storage volume, transferring the primary storage volume scan result to a secondary storage controller for comparison to the secondary CRC scan. Once the primary and secondary CRC scans have been compared the secondary controller may request data stored in the primary storage volume from the primary storage controller to update any non-matching blocks of information on the secondary storage volume. The primary controller may then transfer the appropriate data to the secondary storage controller to allow for the data to be written to the secondary storage volume.

[0011] In a third aspect of the present invention a method for restoring coherency in mirrored storage volumes includes generating a coarse grain bitmap in volatile memory of writing operations occurring on the primary storage volume while communication between the primary and secondary storage volumes is disrupted. Once communications have been reestablished the bitmap may then be utilized to update the secondary storage

volume. In the event of a power surge, reset type error or the like causes an error in the bitmap the method may revert to utilizing CRC scans to reestablish coherency.

[0012] In a further aspect of the invention a method for managing mirrored storage volumes may include conducting a low priority CRC scan of a primary and secondary storage volumes. Once the CRC scans have been conducted the scan results may be compared such that the difference between data blocks may be determine so that data from the primary volume may be updated to the secondary volume.

BRIEF DESCRIPTION OF THE DRAWINGS

[0013] The numerous advantages of the present invention may be better understood by those skilled in the art by reference to the accompanying figures in which:

[0014] FIG. 1 is an overview illustration of an exemplary embodiment wherein a mirrored storage volume system with cyclic redundancy checking (CRC) capability is shown;

[0015] FIG. 2 is an overview illustration of an exemplary embodiment wherein a mirrored storage volume system with CRC capability includes a volatile memory;

[0016] FIG. 3 is a flow chart of an exemplary embodiment of a method of restoring coherency in mirrored storage volumes;

[0017] FIG. 4 is a flow chart of an exemplary method of restoring coherency in mirrored storage volumes including the utilization of a bitmap; and

[0018] FIG. 5 is a flow chart of an exemplary method of utilizing cyclic redundancy checking to ensure the coherency of data in mirrored storage volumes.

DETAILED DESCRIPTION OF THE INVENTION

[0019] Reference will now be made in detail to the presently preferred embodiments of the current invention, examples of which are illustrated in the accompanying drawings.

[0020] Referring generally now to FIGS. 1 through 3, exemplary embodiments of the present invention are shown, wherein a protocol and a controller are suitable for managing the coherency of mirrored data storage volumes.

[0021] Referring now to FIG. 1 a mirrored storage volume system 100 is shown. The mirrored storage volume system 100 of the present invention offers the ability for clients to manage data in an integrated fashion without the drawbacks associated with maintaining a non-volatile bitmap during interruptions in communication. By utilizing the present invention, mirrored storage volumes may effectively manage data storage resources in a cost efficient and effective manner.

[0022] The mirrored storage system 100 may include a primary storage controller 102. The primary storage controller 102 may be linked to a primary storage volume 104 such that the primary storage controller may be capable of managing the reading and writing of data in the primary storage volume 104 as well as performing a block by block cyclic redundancy check (CRC) of the primary storage volume 104. Additionally, the primary storage controller may be capable of comparing the results of various CRC scans to determine differences between scans. For example, the primary storage controller 102 may be capable of comparing the primary storage volume CRC scan to a secondary storage volume 108 CRC scan. During utilization, the primary storage controller 102 may receive input/output operations from clients 112 as well as data transfers from a secondary storage controller 106 linked to a secondary storage volume 108 wherein the secondary storage controller is capable of managing the storage of data in the secondary storage volume 108. Upon the primary storage controller 102 allowing data to be written to the primary storage volume 104, the primary storage controller 102 may transfer the data to be written to the secondary controller 106 so that the secondary storage volume 108 mirrors the primary storage volume 104. Also, the secondary storage controller 106 may be capable of performing a block by block CRC of the secondary storage volume

and comparing the secondary storage volume CRC scan with the CRC scan transferred from the primary storage controller 102.

[0023] The primary storage controller 102 may manage data transfers from the secondary storage controller 106, such that the primary storage controller 102 is allowed to issue read and write requests to the primary storage volume 104, while the secondary storage controller requests permission from the primary storage controller 102 prior to data being transferred to/from the primary storage volume 104. In the present embodiment, the primary storage controller 102 is linked to the secondary storage controller 106 via a communication channel 110.

[0024] The communication channel 110 linking the primary storage controller 102 and the secondary storage controller 106 may be suitable for bi-directional communication. For example, the communication channel 110 may link the primary storage controller 102 in one location with the secondary storage controller 106 located in a different office, geographic location such as a different city or the like.

[0025] During operation of the mirrored storage system 100, an interruption in communication may occur between the primary and secondary controller 102 & 106. As a result, an incoherency may occur between the data stored in the primary and the secondary storage volumes 104 & 108 resulting in the secondary storage volume 108 failing to mirror the contents of the primary storage volume 104.

[0026] In the current embodiment, the mirrored storage system 100, upon reestablishment of the communication channel 110 between the primary and the secondary storage controllers 102 & 106, may be capable of conducting a CRC scan of their respective primary and secondary storage volumes 104 & 108. The primary storage controller 102 may be further capable of transferring the results of the primary storage volume CRC scan to the secondary storage controller 106 for comparison with the

secondary storage volume CRC scan. It is to be understood that the comparison of respective CRC scans may be conducted by the primary storage volume without departing from the spirit and scope of the present invention.

[0027] Once the secondary storage controller 106 has made a comparison of the primary and secondary CRC scans the secondary controller 106 may request the transfer of non-matching blocks of the primary storage volume 104 from the primary controller 102.

[0028] The primary storage controller 102, upon reception of the request for non-matching blocks, transfer the data from the primary storage volume 104 to the secondary storage controller 106 for writing to the secondary storage volume 108.

[0029] In another embodiment of the present invention, the system for managing mirrored storage volumes may be capable of performing CRC scans to determine the coherency of the primary and secondary storage volumes at a set time period such that coherency of the mirrored storage volumes may be ensured.

[0030] In further embodiments of the invention, additional algorithms may be utilized to generate scans for comparison to reestablish coherency. In various embodiments the mirrored storage volume system 100 may conduct at least one of a MD-5 and a SHA-1, based scans to generate a scan for comparison. A mirrored storage volume system may be capable of utilizing scans conducted on the primary and secondary storage volumes to reestablish coherency. It is the intention to include and encompass such changes without departing from the scope and spirit of the present invention.

[0031] Referring now to FIG. 2 in a further embodiment, a mirrored storage system 200 includes a primary storage controller 102 linked to a volatile memory 216 suitable for maintaining a coarse gain bit map. Maintaining a bitmap in non-volatile memory reduces the complexity and cost of mirrored storage volumes systems while including essentially

a redundant CRC system protects against power surges and the like. A coarse grain bitmap may include data representing changes having occurred in a storage volume, such as the primary storage volume. It is to be understood that the granularity of the bitmap may vary depending on the amount of data, the number of write operations having occurred and the like. Upon the occurrence of an interruption in a communication channel 210 between the primary storage controller 202 and a secondary storage controller 206, the primary storage controller 202 creates a coarse grain bitmap indicating the location of changes occurring in a primary storage volume 204 linked to the primary storage controller 202. In this manner the primary storage volume 204 may continue to be written to while there is a break in communication.

[0032] Once the communication channel 210 is reestablished between the primary and secondary storage controllers 202 & 206, the primary controller 202 may utilize the coarse gain bit map to reestablish coherency in the secondary storage volume so that the data in the secondary storage volume 208 mirrors the contents of the primary storage volume 204.

[0033] Should a power failure or the like cause a loss of data in the bitmap the primary storage controller 202 reverts to a CRC scan of the primary storage volume 204, once the communication channel 210 is reestablished, and transmitting a message to the secondary storage controller 206 to initiate a CRC scan of the secondary storage volume 208. The primary and secondary CRC scans are compared by the secondary storage controller 206 and non-matching blocks requested from the primary controller 202 for writing to the secondary storage volume 208.

[0034] In reference to FIG. 3, an exemplary method for establishing coherency in mirrored storage volumes 300 is shown, wherein coherency of mirrored storage volumes is restored. Initially, upon the reestablishment of a communication channel 302 after an interruption, primary and secondary storage controllers conduct a block-by-block CRC

scan of their respectively linked primary and secondary storage volumes. Conducting a CRC scan 304 includes the utilization of a algorithm to calculate a binary code as a result of arithmetic operations on the data included in a storage volume, such as the primary storage volume in the case of the primary storage controller.

[0035] After the primary and the secondary storage controllers conduct the CRC scans 304, the result of the primary CRC scan may be transferred to the secondary storage controller 306. Transferring the primary scan result 306 may include transmitting the data generated as a result of the CRC scan to the secondary controller through a communication channel such as the communication channel 110 in FIG. 1.

[0036] Once the secondary storage controller receives the result of the primary scan, the secondary controller compares the result of the secondary storage volume CRC scan result to that of the primary storage volume CRC scan result. Comparing the primary CRC scan 308 to the secondary storage volume scan may include comparing scan results for blocks of data in the primary storage volume with the corresponding CRC scan for corresponding blocks of data in the secondary storage volume.

[0037] Following the comparison of the CRC scans for the primary and secondary storage volumes 308, the secondary storage controller may request the transfer of non-matching blocks from the primary storage volume via the primary storage controller. Requesting non-matching blocks 310 includes data blocks where the secondary storage volume fails to mirror the primary storage volume.

[0038] The primary storage controller may transfer the non-matching blocks of data by initiating a reading of the primary storage volume upon the receipt of the request for the non-matching blocks 310. Transmitting non-matching blocks 312 by the primary storage controller may include initiating a reading of the non-matching blocks in the primary storage volume. The data from the non-matching blocks of the primary storage volume

are transmitted to the secondary storage controller via the communication channel between the primary and secondary controller.

[0039] Upon reception of the data from the primary storage volume, the secondary controller writes the data transferred from the primary storage volume to the secondary storage volume. Writing the non-matching blocks 314 may include writing information to the secondary storage volume such that the non-matching block on the secondary volume is updated with the primary storage volume data from the corresponding block.

[0040] Furthermore, in additional embodiments the method may utilize various algorithms to restore coherency. The method for restoring coherency may employ additional algorithms such as at least one of a MD-5 and a SHA-1 scan to generate scans for comparison. It is the intention to include and encompass such changes without departing from the scope and spirit of the present invention.

[0041] In a further embodiment, with reference to FIG. 4, a method of restoring coherency in mirrored storage volumes 400 may include utilizing a coarse grain bitmap maintained in a volatile memory linked to a primary storage controller. A coarse gain bitmap may include data representing changes having occurred in a storage volume, such as the primary storage volume. It is to be understood that the granularity of the bitmap may vary depending on the amount of data, the number of write operations having occurred and the like. The generation and utilization of a coarse grain bitmap stored in volatile memory may allow a mirrored storage volume system to maintain the changes without the need for maintaining the bitmap in non-volatile memory. Thus, the need for costly hardware is reduced while retaining the capability of reverting to a CRC method of establishing coherency should a power failure, reset type error or the like occur and corrupt the bitmap.

[0042] In the present embodiment, upon an interruption in communication 402 the primary storage controller may generate a coarse grain bitmap of changes in the primary storage volume. Generating a bitmap 404 may include information suitable for determining data changes having occurred while a communication channel is disrupted.

[0043] Once the communication channel between the primary storage controller and a secondary storage controller has been reestablished 406, the primary storage controller may utilize the bitmap 406 to update changes to the secondary storage volume. Should a power surge, reset type error or the like occur, the method of restoring coherency in mirrored storage volumes 400 may revert to the CRC approach method 408 for reestablishing coherency as set forth above, with specific reference to the method of restoring coherency in mirrored storage volumes 300 of FIG. 3.

[0044] Referring to FIG. 5, in an additional embodiment a method for managing mirrored storage volumes 500 is shown wherein the coherency of mirrored storage volumes is safeguarded against hidden corruption, due to faulty storage media or the like. Initially, at a set time period 502, a primary and a secondary storage controllers may conduct a low-priority block by block CRC scan of their respectively linked primary and secondary storage volumes. Conducting a CRC scan 504 may include the utilizing a algorithm to calculate a binary code as a result of arithmetic operations on the data included in a storage volume, such as the primary storage volume in the case of the primary storage controller.

[0045] After the primary and the secondary storage controllers conduct the CRC scans 504, the result of the primary CRC scan may be transferred to the secondary storage controller 506. Transferring the primary scan result 506 may include transmitting the data generated as a result of the CRC scan to the secondary controller through a communication channel such as the communication channel 110 in FIG. 1.

[0046] Once the secondary storage controller receives the result of the primary scan, the secondary controller compares the result of the secondary storage volume CRC scan result to that of the primary storage volume CRC scan result. For instance the primary CRC scan 508 to the secondary storage volume scan may compare scan results for blocks of data in the primary storage volume with a corresponding CRC scan for corresponding blocks of data in the secondary storage volume.

[0047] Following the comparison of the CRC scans for the primary and secondary storage volumes 508, the secondary storage controller may request the transfer of non-matching blocks from the primary storage volume via the primary storage controller. Requesting non-matching blocks 510 includes data blocks where the secondary storage volume fails to mirror the primary storage volume.

[0048] The primary storage controller may transfer the non-matching blocks of data by initiating a reading of the primary storage volume upon the receipt of the request for the non-matching blocks 510. Transmitting non-matching blocks 512 by the primary storage controller may include initiating a reading of the non-matching blocks in the primary storage volume and transmitting the data from the non-matching blocks of the primary storage volume to the secondary storage controller via the communication channel between the primary and secondary controller.

[0049] Furthermore, upon the reception of the data from the primary storage volume the secondary controller may write the data transferred from the primary storage volume to the secondary storage volume. Writing the non-matching blocks 514 may include writing information to the secondary storage volume such that the non-matching block on the secondary volume may be updated with the primary storage volume data from the corresponding block.

[0050] Furthermore, in additional embodiments the method may utilize various algorithms to restore coherency. The method for restoring coherency may employ additional algorithms such as at least one of a MD-5 and a SHA-1 to generate scans for comparison. It is the intention to include and encompass such changes without departing from the scope and spirit of the present invention.

[0001] It is believed that the CYCLIC REDUNDANCY CHECKING FOR MANAGING THE COHERENCY OF MIRRORED STORAGE VOLUMES of the present invention and many of its attendant advantages will be understood by the forgoing description. It is also believed that it will be apparent that various changes may be made in the form, construction and arrangement of the components/ steps thereof without departing from the scope and spirit of the invention or without sacrificing all of its material advantages. The form herein before described being merely an explanatory embodiment thereof. It is the intention of the following claims to encompass and include such changes.